



Carnegie Mellon University
Computer Science Department
5000 Forbes Avenue
Pittsburgh, PA 15213-3891 USA

April 14, 1997

ESC/AXS Harry Koch
ARPA Agent
5 Eglin Street
Building 1704, Room 205
Hanscom AFB, MA 01731-2116

Dear Harry:

RE: Contract F19628-95-C-0050
"The Fox Project: Advanced Languages for Systems Software"
#1-52220

Enclosed is the quarterly R&D Status Report covering our research progress during the period January 1 through March 31, 1997. Should you have any questions, please do not hesitate to contact me at 412/268-3853.

Have a wonderful day!

Best regards,

A handwritten signature in cursive script that reads "Rosie".

Rosie Hornyak

/ r m h
Enclosures

Copy to: G. Koob, DARPA/ITO
C. Stephan, ESC/AXK
DARPA Technical Library
Office of Naval Research
Defense Technical Information Center/OCC

P. Lee, CMU
R. Harper, CMU
M. Brendel, CMU
A. Stoltzfus, CMU
E. Biagioni, CMU

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

19970421 203

The Fox Project: Advanced Development of Systems Software

R&D Status Report
January 1 to March 31, 1997

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

This research is sponsored by the Defense Advanced Research Projects Agency, DoD, through ARPA Order 8313, and monitored by ESD/AVS under contract F19628-95-C-0050. Views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the United States Government.

The long-term objectives of the Carnegie Mellon Fox Project are to improve the design and construction of systems software and to further the development of advanced programming languages. We use an advanced programming language in the design and construction of systems software, including operating systems, network protocols, and distributed systems. The language is based on Standard ML (SML), a modern functional programming language that provides polymorphism, first-class functions, exception handling, garbage collection, a parameterized module system, static typing, and a formal semantics. This Project involves several faculty members and spans a wide range of research areas, from (1) experimental development of systems software to (2) advanced compiler development to (3) language design.

1 Research Progress

For each of the three areas listed above, we report on the research accomplishments during the first calendar quarter of 1997, and the research objectives for the second quarter of 1997.

1.1 Experimental Development of Systems Software

Accomplishments (January-March):

- Improved the performance of the FoxNet. TCP (without checksums) now runs at 6 Megabits/second on a 10 Megabits/second Ethernet.
- Studied the sources of inefficiency in the FoxNet. These included performance bugs in TCP, the excessive use of abstraction in the marshaling and unmarshaling code, and inefficient implementation of word access and word shifting operations in the SML/NJ compiler.
- Developed a flexible set of performance analysis tools for measuring different parts of the FoxNet.
- Investigated poor code generated by the SML/NJ compiler in the handling of byte arrays and bit shift commands.
- Developed a FoxNet device protocol that supports an ATM network, and integrated it into the FoxNet to support existing applications.

Objectives (April-June):

- Identify the storage leak currently observed in the FoxNet web server.
- Fix the SML/NJ compiler to improve byte array handling, and bit shift functions, to improve the marshaling and unmarshaling performance of the FoxNet.
- Revise the "Safe-for-Space Threads in Standard ML" for journal publication.

Noteworthy publications:

- "Safe-for-Space Threads in Standard ML", by Edoardo Biagioni, Ken Cline, Peter Lee, Chris Okasaki, and Chris Stone. 1997 ACM SIGPLAN Workshop on Continuations, Paris, January 14, 1997.

1.2 Language Design

Accomplishments (January-March):

- Designed the "phase-splitting transformation" that reduces the module constructs of Standard ML to a much simpler polymorphic lambda calculus. This will allow us to extend the TIL compiler to handle the SML module system.
- Developed a working prototype of sort analysis, to improve the type analysis of SML programs.
- Designed three prototype "certifying compilers" for automatically generating proof-carrying code target programs from high-level source programs.
- Proved the soundness of proof representation and validation in Proof-Carrying Code.

Objectives (April-June):

- Complete revision of the Milner Festschrift paper described below.
- Complete prototype of new proof checker for PCC using explicit substitutions.
- Start the implementation of a Certifying Compiler prototype.

Noteworthy publications:

- "A Type-Theoretic Interpretation of Standard ML 1996", by Robert Harper and Chris Stone. Accepted for publication (subject to revisions) for the Milner Festschrift to be published at the end of this year. This paper describes in detail the type-theoretic interpretation of Standard ML, a critical ingredient for building a type-based compiler for the language.
- "Subtyping is not a good 'Match' for object-oriented languages", by Kim B. Bruce, Adrian Fiech, and Leaf Petersen. To appear in Proceedings of ECOOP '97
- "Catenable Double-Ended Queues", by Chris Okasaki. To appear in ICFP'97.
- "Three Algorithms on Braun Trees", by Chris Okasaki. To appear in the Journal of Functional Programming.
- "Proof-Carrying Code", by George Necula, POPL '97.

- "Even Higher-Order Functions for Parsing or Why Would Anyone Ever Want To Use a Sixth-Order Function?", by Chris Okasaki. To appear in the Journal of Functional Programming.
- "Implementing Bit-addressing with Specialization", by Scott Draves. To appear in ICFP97.

1.3 SML Compiler and System Development

Accomplishments (January-March):

- Implemented propagating of type information through MLRisc back-end as part of TIL compiler project.
- Completed the implementation of the elaboration phase of the TILT compiler for full Standard ML. This is a critical step towards extending TIL to the full SML language.

Objectives (April-June):

- Finish initial implementation of TIL/ML compiler.
- Finish integrating MLRisc back-end into TIL compiler.
- Complete a first version of the TILT compiler for full Standard ML (in conjunction with Greg Morrisett at Cornell).
- Complete working sort checker for ML core.
- Add support for lazy evaluation to TIL.

2 Capital Equipment Purchases

- 1 220-0320 Dell 6200/OP GXpro, w/Pentium Pro 200Mhz/256K, Integrated, \$6,467.00
- 1 220-0320 Dell 6200/OP GXpro w/Pentium Pro 200Mhz/256K, Integrated, \$5,681.00
- 2 ST1515ON 2 4GB 3.5" Narrow SCSI-2 Disks, x2C13520 in 3.5" Dual Bay Enclosure, \$3,850.00
- 1 ST1515ON 1 4GB 3.5" Narrow SCSI-2 Disks, C13520 in 3.5" Dual Bay Enclosure, \$1,030.00

3 Key Personnel Changes

- Peter Lee was appointed to a two-year term on the Army Science Board. He is currently participating on an Army Science Board special study panel on Distance Learning.

4 Noteworthy Meetings

- POPL'97 (Paris, Jan 14-17)
Peter Lee was the Co-General Chair for the conference. Also, he presented a paper on "Safe-for-Space Threads in Standard ML", at the 1997 Continuations Workshop, co-located with POPL'97.
- NRC panel on security infrastructure (Irvine, CA, Feb.5-6).
Peter Lee was invited to participate in this panel meeting, to present work on Proof-Carrying Code, and discuss its possible applications to computer security related issues.
- Distinguished Lecture, Univ. of Illinois (Mar.10).
Peter Lee gave a distinguished lecture on the topic, "Proofs, Types, and Mobile Code".
- DARPA Workshop on Foundations for Secure Mobile Code (Monterey, CA, Mar.26-28).
Peter Lee and George Necula attended this workshop. Necula presented work on Proof-Carrying Code. Peter Lee and Joan Feigenbaum (AT&T Labs) presented a joint position paper on integrating Proof-Carrying Code into a Trust Management infrastructure.

5 Administrative Data

```
><ADMINISTRATIVE_AND_FINANCIAL_DATA>
>
><DATE_PREPARED>
10 APR 97
></DATE_PREPARED>
>
><AWARD_NUMBER>
F19628-95-C-0050
></AWARD_NUMBER>
>
><AWARD_AGENT>
Hanscom Air Force Base
></AWARD_AGENT>
>
><AWARD_TITLE>
The Fox Project: Advanced Languages for Systems Software
></AWARD_TITLE>
```

```
>
><ACTUAL_START_DATE>
08 JUN 95
></ACTUAL_START_DATE>
>
><ACTUAL_START_DATE_COMMENT>
>
></ACTUAL_START_DATE_COMMENT>
>
><OFFICIAL_AWARD_END_DATE>
30 JUN 97
></OFFICIAL_AWARD_END_DATE>
>
>
><CURRENT_AWARD_FUNDING_PROFILE_BY_YEAR>
>
>FY95 Base 922,250
>FY96 Base 0; Option1 964,201
>FY97 Base 0; Option1 0; Option2 1,008,341; Option3 423,046
>FY 98 Base 0; Option1 0; Option2 0; Option3 0; Option4
444,260
>TOTAL Base 922,250; Option1 964,201; Option2 1,008,341; Option3 423,046;
Option4 444,260
>
></CURRENT_AWARD_FUNDING_PROFILE_BY_YEAR>
>
><TOTAL_FUNDS_PROVIDED_TO_DATE>
>
>Base 922,250
>Option1 964,201
>
></TOTAL_FUNDS_PROVIDED_TO_DATE>
>
><ACTUAL_FUNDS_EXPENDED_TO_DATE>
>
>Base 922,250
>Option1 638,297
>
></ACTUAL_FUNDS_EXPENDED_TO_DATE>
>
><DATE_CURRENT_FUNDING_EXPENDED>
>30 JUN 97
></DATE_CURRENT_FUNDING_EXPENDED>
>
><FUNDING_BY_FISCAL_QUARTER>
>
>Option 1 PLANNED PROVIDED DELTA
>Funding for 7/96-9/96 276,227 276,227 0
>Funding for 10/96-12/96 276,227 276,227 0
> FY96 Option 1 552,454 552,454 0
>
>Funding for 1/97 - 3/97 205,873 205,873 0
>Funding for 4/97 - 6/97 205,874 205,874 0
```

> FY97 Option 1	411,747	411,747	0
>			
>Option 2	PLANNED	PROVIDED	DELTA
>Funding for 7/97 - 9/97	252,085	0	252,085
>Funding for 10/97 - 11/97	168,057	0	168,057
> FY97 Option 2	420,142	0	420,142
>			
>Funding for 12/1/97- 12/31/97	84,028	0	84,028
>Funding for 1/98 - 3/98	252,085	0	252,085
>Funding for 4/98 - 6/98	252,086	0	252,086
> FY98 Option 2	588,199	0	588,199
>			
Total Option 2	1,008,341	0	1,008,341

>Option 3	PLANNED	PROVIDED	DELTA
>Funding for 7/96 - 9/96	105,761	0	105,761
>Funding for 10/96 - 12/96	105,762	0	105,762
> FY96 Option 3	211,523	0	211,523
>			
>Funding for 1/99 - 3/99	105,761	0	105,761
>Funding for 4/99 - 6/99	105,762	0	105,762
> FY97 Option 3	211,523	0	211,523
>			
>Option 4	PLANNED	PROVIDED	DELTA
>Funding for 7/97 - 9/97	111,065	0	111,065
>Funding for 10/97 - 12/97	111,065	0	111,065
> FY97 Option 4	222,130	0	222,130
>			
>Funding for 1/98 - 3/98	111,065	0	111,065
>Funding for 4/98 - 6/98	111,065	0	111,065
> FY98 Option 4	222,130	0	222,130

```

></FUNDING_BY_FISCAL_QUARTER>
>
><DATE_OF_FINANCIAL_DATA>
>31 MAR 97
></DATE_OF_FINANCIAL_DATA>
>
><ANYTHING_ELSE_YOU_NEED>
>Option 3 would have supplemented this year's option 1 money; option 4 would
>supplement next year's option 2 money.
></ANYTHING_ELSE_YOU_NEED>
>
></ADMINISTRATIVE_AND_FINANCIAL_DATA>
>

```